

REMARKSI. Introduction

In response to the Office Action dated January 16, 2003, claims 37, 50, and 63 have been cancelled and claims 36, 49, and 62 have been amended. Re-examination and re-consideration of the application, as amended, is requested.

II. Claim Amendments

Applicants' attorney has made amendments to the claims as indicated above. These amendments were made solely for the purpose of clarifying the language of the claims, and were not required for purposes of patentability.

III. Office Action Non-Art Rejections

In paragraph 1, the Office Action rejects claims 38, 40, 51, 53, 64, and 66 under 35 U.S.C. § 112 as being indefinite for failing to particularly point out in distinctly claim the subject matter which the applicant regards as the invention.

According to the Office Action, the term "operator" may designate a person or an abstract operator, and that the language in these claims should be modified to distinguish between these two possible meanings of the word "operator".

The Applicants respectfully traverse these rejections, because the meaning of the term "operator" is apparent from both the claim itself and the specification. For example, claim 1 recites and "operator tree", a phrase which certainly cannot be reasonably confused with a person. The Applicants' specification also includes numerous references to the term "operator," and in none of these contexts, can the word "operator" be reasonably confused with a "person."

IV. The Cited References and the Subject InventionA. The Callahan Reference

U.S. Patent No. 6,230,313, issued May 8, 2001 to Callahan et al. discloses a system for conducting performance analysis for executing tasks. The analysis involves generating a variety of trace information related to performance measures, including parallelism-related information, during

C

execution of the task. In order to generate the trace information, target source code of interest is compiled in such a manner that executing the resulting executable code will generate execution trace information composed of a series of events. Each event stores trace information related to a variety of performance measures for the one or more processors and protection domains used. After the execution trace information has been generated, the system can use that trace information and a trace information description file to produce useful performance measure information. The trace information description file contains information that describes the types of execution events as well as the structure of the stored information. The system uses the trace information description file to organize the information in the trace information file, extracts a variety of types of performance measure information from the organized trace information, and formats the extracted information for display. The system can use default or user-defined functions to extract and format trace information for display. After the system displays one or more types of performance measure information, a user of the system can then interact with the system in a variety of ways to obtain other useful performance analysis information.

#### B. The Epperson Reference

U.S. Patent No. 5,754,771, issued May 19, 1998 to Epperson et al. discloses a maximum receive capacity specifying query processing client/server system replying up to the capacity and sending the remainder upon subsequent request. An Interactive Television (ITV) Client/Server system comprises one or more Clients (e.g., set top boxes or "STBs") connected to a Server module, via a Broadband Communication Network, is described. The system provides a generic mechanism for deciding how much or how little data is to be sent in response to a request from a Client. Specifically, each Client is allowed to describe ahead of time how much memory (approximately or exactly) it is allowing for the query result. This is communicated by the Client as part of the query itself. Regardless of what other activity the Client undertakes, the server maintains the context of the original query, until the Client terminates the connection. In a similar manner, when other simultaneous queries are open, the context is maintained for each query until the connection is dropped (or the query terminates). The Client can ask for more data (i.e., rows) on any open query. To optimize operation of the environment, requests from the Clients (i.e., queries) are de-multiplexed down to a set of worker threads, available at the server, which carry out the actual work.

of query processing. Because of this approach, a request for more data from a Client (i.e., set top box) can "land" on any worker thread. This gives queries the ability to "jump" from one thread of execution to another. In this manner, a finite number of threads may be employed to service a large query pool, with maximum throughput.

#### C. The Posse Reference

U.S. Patent No. 5,544,175, issued August 6, 1996 to Posse discloses a method and apparatus for the capturing and characterization of high-speed digital information. A digital signal detector for sampling the state of a high speed digital signal occurring at a test node in a digital circuit which exhibits the same behavior with repeated applications of the same inputs. The digital signal detector samples the state of a test circuit node at discrete intervals in time and stores the digital levels with timing reference information. The stored information is then compared to the expected behavior of the tested node for analysis of delay-type and other parametric faults, such as performance faults. The digital signal detector includes a state discriminator which determines the state of the input digital signal by comparing its voltage level to one or more threshold voltages. The digital signal detector also includes a memory unit which stores the signal state information upon receipt of a memory strobe signal generated by a memory delay control circuit which delays the application of the memory strobe signal by a predetermined amount of time after the receipt of a clock pulse. The stored digital signal states can be later retrieved from the memory unit for evaluation of the input digital signal.

#### D. The Bhargava Reference

U.S. Patent No. 5,680,603, issued October 21, 1997 to Bhargava et al. discloses a method and apparatus for reordering complex SQL queries containing inner and outer join operations. A method and apparatus for reordering complex SQL queries containing joins, outer and full outer joins. The method and apparatus first translates the query into a hypergraph representation. Required sets, conflict sets and preserved sets are then generated for the query hypergraph. Using the required sets, a plurality of plans are enumerated, wherein the plans represent associative reorderings of relations in the query. SQL operators are selectively assigned to each of the enumerated plans using the conflict sets and/or preserved sets, so that the results from the plans are

identical to the original query. A novel Modified General Outer Join (MGOJ) operator may be assigned to the root of a sub-tree, wherein the MGOJ operator is a compensation operator. The operator assignment is performed recursively for the root of each sub-tree in the plan. One of the enumerated plans (generally the most optimal) is then selected for execution.

#### E. The Kimmerly Reference

U.S. Patent No. 5,628,017, issued May 6, 1997 to Kimmerly et al. discloses a method and system for providing event-response capabilities to pseudocode. A method and system for providing event-response and monitoring capabilities to a pseudocode program operating in a message or event-based operating environment. The method and system allow the pseudocode program to asynchronously respond to events via one or more trap routines located in the pseudocode program. The method and system also allow the pseudocode program to synchronously monitor events and receive parameters from operating system routines via callback routines located in the pseudocode program. The system includes an execution engine for executing instructions of the pseudocode program. An event-response routine in the pseudocode program includes instructions for responding to the event. An event-response dispatcher is adapted to receive information identifying the event-response routine, save the execution state of the execution engine, cause the execution engine to execute the event-response routine, and restore the execution state of the execution engine that existed before the event-response routine was executed. A transfer routine accessible in response to a call triggered by the event accesses the event-response dispatcher, passes information identifying the event-response routine to the event-response dispatcher, and returns control to the execution engine after the event-response dispatcher restores the state of the execution engine and returns to the transfer routine.

#### F. The Rhodes Reference

U.S. Patent No. 6,073,110, issued June 6, 2000 to Rhodes et al. discloses an activity based equipment scheduling method and system. A computer based equipment scheduling system uses activity definition data to schedule equipment. The method may be carried out by a networked computer system that receives activity definition data representing an activity, such as a volleyball game or board meeting, that occurs in a building. This data is stored in a database for use by



multiple nodes. The activity definition data may include text or graphical data indicating an activity name, data representing a list of zones that the activity will affect and data representing the mode of possible operation of each building zone affected by the activity. The method also includes assigning zone mode data to the activity data for use in controlling building resources and then automatically controlling the building resources to attain a zone mode of operation assigned to the activity definition data.

#### G. The Junkin Reference

U.S. Patent No. 6,493,717, issued December 10, 2002 to Junkin discloses a system and method for managing database information to be presented in HTML format for retrieval and display by a Web browser. Database information is managed by responding to a user selection delivered by browser software by retrieving the contents of portions of a database and constructing an HTML-compatible presentation of the contents in accordance with definitions that hierarchically link the portions separately from any database-provided links between the portions.

#### H. The Harel Reference

U.S. Patent No. 5,873,081, issued February 16, 1999 to Harel discloses a method and mechanism for filtering incoming documents against user queries. A plurality of user queries including terms connected by logical operators is received. Terms and sub-expressions are combined into distinct sub-expressions and embedded into a directed acyclic graph (DAG) having a plurality of nodes. Each node in the DAG includes pointers to any successor nodes thereof, the terms in the queries are embedded as source nodes in the graph, and the operators embedded as internal nodes. When a document is received, the document is evaluated against the nodes in the DAG by comparing the relevant terms in the document with the source nodes in the DAG representative thereof. For each term that matches a source node, the internal successor node of the matched source node is evaluated based on the logical operator represented by the successor node and truth information of the predecessor nodes thereto, thereby determining a truth value of the internal successor node. Information is returned indicative of which of the successor nodes were evaluated as true. From that information, the queries which matched the document and the users corresponding thereto can be determined.



### I. The Zhou Reference

U.S. Patent No. 5,995,511, issued November 30, 1999 to Zhou et al. discloses a queue control system for use in connection with the transfer of information, in the form of information transfer units, in a digital network. The network provides a plurality of service rate classes, based on, for example transmission rates for the various paths. The information buffer control subsystem includes a information transfer unit receiver, a information transfer unit buffer and a group controller. The information transfer unit receiver receives the information transfer units, and the buffer is provided to buffer the received information transfer units prior to transmission. The group controller controls the buffering of information transfer units received by the information transfer unit receiver in the buffer. In that operation, the group controller aggregates the information transfer units for each path in the buffer according to respective service rate classes, in particular aggregating the information transfer units for each path in a queue and further aggregating the queues for the paths associated with each service rate class in a queue. A transmission scheduler is also disclosed for use in transferring information, in the form of information transfer units, each associated with a path, in a digital network. The network provides a plurality of service rate classes, based on, for example, transmission rates for the various paths. The information transfer units for each path in a path queue, and the path queues for the paths associated with each service rate class are aggregated in a service rate queue. The transmission scheduler includes a information transfer unit selector for selecting from among the service rate queues, one path queue to provide a information transfer unit for transmission, and a information transfer unit transmitter for transmitting the information transfer unit provided by the selected path queue.

### J. The Bamford Reference

U.S. Patent No. 6,243,702, issued June 5, 2001 to Bamford et al. discloses a method and system for removing propagation delays between a plurality of database servers that have access to a common database is provided. According to the method, each database server is associated with a logical clock. In response to initiating a commit of a transaction executing on a database server, a commit time for the transaction is determined and broadcast to one or more other database servers. The broadcast is overlapped with a transaction log force. Upon receiving the commit time, the

database servers compare the transmitted commit time to the time indicated by their logical clock. If the commit time is greater than the time indicated by their logical clock, the database server sets its logical time to reflect a time that is at least as recent as the time reflected by the transmitted commit time.

#### K. The Naidu Reference

U.S. Patent No. 5,752,002, issued May 12, 1998 to Naidu et al. discloses a method and apparatus for optimizing performance of a computer system component design by performance analysis of a simulation of the design. The method of the present invention comprises providing the computer system component design to an analyzing apparatus and carrying out a simulation run of the design. During the simulation run, operation data is generated cycle by cycle, and the generated operation data is collected and stored in a log file. The log file is input to a parser and the operation data is sequentially parsed to produce parsed data. Statistical calculations are then performed on the parsed data, and the performance results are output to the designer in graphical form. The performance information can be used to enhance performance of the computer system component prior to its fabrication.

#### L. The Filepp Reference

U.S. Patent No. 5,594,910, issued January 14, 1997 to Filepp et al. discloses a distributed processing, interactive computer network and method of operation. The network is designed to provide very large numbers of simultaneous users access to large numbers of applications which feature interactive text/graphic sessions. The network includes one or more host computers having application data stores; a plurality of concentrator computers, also including application data stores, the concentrator computers being connected in groups of one or more to each of the host computers; and a plurality of reception system computers connected in groups of one or more to each of the concentrator computers, the reception system computers being arranged so that respective users can request interactive applications at the reception system computers. In accordance with the design, the reception system computers also include application data stores. The method for operating the network includes steps for generating the interactive text/graphic sessions from objects that include data and/or program instructions. Additionally, the method features steps



for distributing objects among the data stores of the network computers, and, thereafter, permitting the reception system computer at which an application is requested to selectively collect objects required for the application from the network and the respective reception system so that the requested application may be presented at the reception system based on the objects collected. This operation decreases processing demand on the higher-level network elements, permitting them to function primarily as data supply and maintenance resources, thereby reducing network complexity, cost and response time.

#### **M. The Carey Reference**

U.S. Patent No. 6,285,997, issued September 4, 2001 to Carey et al. discloses a method, apparatus, and article of manufacture for a computer-implemented technique for query optimization with deferred updates and autonomous sources. An object-oriented query is executed to retrieve data from a database. The database is stored on a data storage device connected to a computer. The object-oriented query is transformed into subqueries, wherein at least one subquery is directed against a database, and wherein one subquery is directed against an object cache. Each subquery that is directed against a database is executed to retrieve data from the database into the object cache. The subquery that is directed against the object cache is executed to retrieve data for the query, wherein the data incorporates updates to the object cache and updates to the database.

#### **V. Office Action Prior Art Rejections**

In paragraph (3)-(4), the Office Action rejected claims 36, 43, 45, 49, 56, 58, 62, 69 and 71 under 35 U.S.C. § 103(a) as unpatentable over Callahan, II et al., U.S. Patent No. 6,230,313 (Callahan), Epperson et al., U.S. Patent No. 5,754,771 (Epperson), and Posse, U.S. Patent No. 5,544,175 (Posse). The Applicants respectfully traverse these rejections.

With respect to the claim 36, 49, and 62: claim 36 has been amended to recite the features of claim 37. As amended, claim 36 recites:

*A method of monitoring an execution of a query performed by a database system having a query coordinator and at least one data server, wherein the query execution comprises at least one execution thread, the method comprising the steps of:*

*for each thread, generating first execution trace information in the query coordinator, wherein the first execution trace information comprises an execution plan in terms of one or more operator trees;*

C

*for each thread, generating second execution trace information in the data server; and writing the first execution trace information and the second execution trace information to at least one execution log file.*

According to the Office Action, the Callahan referenced discloses "... for each thread, generating first execution trace information ..." and " for each thread, generating second execution trace information..." as follows:

task. In particular, target source code of interest is compiled 20 in such a manner that executing the resulting target executable code will generate execution trace information related to a variety of performance measures for the one or more processors and protection domains used. The compiler will

and

be parallelized. For example, a region of code that loops 25 multiple times with sequential values of a variable (e.g., 'loop for x from 1 to 10') may be able to be divided among multiple processors so that each instance of the loop (with a different value of the variable) can be executed by a different thread simultaneously. After the compiler has identified

As the Office Action acknowledges, Callahan, Epperson, and Posse do not teach generating *first execution trace information comprising an execution plan in terms of one or more operator trees.*

Referring to the rejection of claim 37, the Office Action indicates that the Bhargava reference teaches the use of execution plans and the use of operator trees as follows:

FIGS. 10A, 10B, and 10C together are a flowchart illustrating the method of performing conflict-free operator assignments for the association trees according to the present invention; and 25

and;

present invention. The conflict-free operator assignment 55 attempts to generate an operator tree for a given association tree by assigning joins, outer joins and full outer joins to the interior nodes of the association tree. If it is not possible to

The Applicants agree that the foregoing teach the use of execution plans and teaches operator trees. But the foregoing does not suggest generating execution trace information in the

form of operator trees. The Office Action argues that this substitution would be obvious to one of ordinary skill in the art because Callahan uses "tasks" and Bhargava "use[s] nodes."

To be precise, Callahan generates raw trace information, that is later modified by the trace information display or (TID) for presentation to the user. The raw trace information includes information such as that which is illustrated in FIG. 4. In Callahan, each thread does not report trace information in the form of "tasks", but rather, as *raw* information that is later assembled and formatted into a form is desired by the user.

Callahan uses "tasks" but clearly, does not report trace information in the form of "tasks." Callahan therefore either (1) teaches that there is no relationship between the "task" organization and the form of the reported execution trace information, or (2) teaches that the execution trace information should be reported *differently*. In either case, Callahan can only be said to teach away from claim 1.

Given the foregoing, it does not follow that if the Bhargava, which the Office Action indicates "use[s] nodes", would obviously report trace information the form of execution plans having one or more operator trees.

"A reference may be said to teach away when a person of ordinary skill, upon reading the reference, would be discouraged from following the path set out in the reference, or would be led in a direction divergent from the path that was taken by the Applicants. The degree of teaching away will of course depend on the particular facts; in general, a reference's disclosure will teach away if it suggests that the line of development flowing from the reference's disclosure is unlikely to be productive of the result sought by the Applicants. *In re Gurley*, 27 F.3d 551, 553, 31 U.S.P.Q.2d 1130 (Fed. Cir. 1994).

Claim 1 recites that the execution trace records are generated in separate entities (the data server and the query coordinator). The Applicants can ascertain nothing in the Callahan, Epperson, or Posse references that disclose this feature.

Finally, the Office Action acknowledges that the Callahan and Epperson references do not teach writing execution trace information to at least one execution log file. However, the Office Action alleges that the following section of the posse reference nonetheless teaches this feature:

At each successive execution of the test program, the <sup>25</sup> value of the low-order address bits is incremented by one. This advances the sampling of the digital signal (which is regenerated with each test program execution) in discrete intervals. Referring to the above example, during the next

and,

In a preferred embodiment of the present invention, tester 500 interfaces with a simulator 508. Simulator 508 provides reference data representing the expected behavior of test node 512 to control unit 502 via reference data line 507.   
50 Typically, the simulation of digital circuit 515 is performed prior to (and independently of) the testing of digital circuit 513. The results of the simulation are then provided to control unit 502 in what is referred to as a simulation log file. The simulation log file contains threshold crossings and timing reference data for the test node 512 based upon the application of a given set of test inputs 513 to digital circuit 515. In the preferred embodiment of the present invention,

The foregoing refers to a simulator 508 which interfaces with a circuit tester 500. Results from the simulation are provided to a control unit 502 in what is referred to as a simulation log file, which has threshold crossings and timing reference data for test nodes based upon the application of a given set of test inputs to a digital circuit. The Applicants respectfully point out that this reference refers to an entirely different art than that of the Applicants' invention, and refers to collecting *digital circuit simulation* results, not execution trace information with operator trees. The applicants respectfully disagree that one of ordinary skill in the art would be motivated to alter the teachings off Callahan, Epperson, and Bhargava as described in an unrelated reference in an entirely different art (digital circuit simulation and analysis).

Claims 43, 49, 56, 62, and 69 include limitations analogous to those of claim 36, and are patentable for the same reasons.

Regarding Claims 45, 58, and 71: According to the Office Action, the step of reconstructing the execution trace information from the log file is disclosed in the Posse reference. However,

- column 8, lines 28-33 of the Posse reference refers to reconstructing a test period of an input signal;
- column 9, lines 27-29 of the Posse reference refers to testing of a digital signal;

- column 10, lines 1-3 of the Posse reference describes engineering models that generate library files containing threshold crossing and “dependency” and “condition information”; and
- column 10, lines 54-57 of the Posse reference disclose reconstructing simulation data.

None of the foregoing discloses reconstructing execution trace information with operator trees from an execution log file.

In paragraph (11), the Office Action rejected claims 38, 51, and 64 under 35 U.S.C. §103(a) as being unpatentable over Callahan, Epperson, Posse, and Bhargava as applied to claims 37, 50, and 63 above respectively, and further in view of Kimmerly et al., U.S. Patent No. 5,628,017 (Kimmerly) and Rhodes et al., U.S. Patent No. 6,073,110 (Rhodes). The Applicants respectfully traverse these rejections.

Claim 38 recites that the execution trace information further comprises operator dispatch information. The Office Action suggests that Kimmerly discloses the use of operator dispatch information at column 3, lines 9-16. However, Kimmerly does not disclose including operator dispatch information in execution trace information, as claim 38 recites. Further, the Office Action provides no motivation for using such operator dispatch information, only that it is known to be used in other contexts. Accordingly, the Applicants do not believe the Office Action has presented a *prima facie* case of obviousness under 35 U.S.C. § 103.

Claims 51, and 64 include limitations analogous to those of claim 38, and are patentable for the same reasons.

In paragraph (12), the Office Action rejected claims 39, 52, and 65 under 35 U.S.C. §103(a) as being unpatentable over Callahan, Epperson, and Posse as applied to claims 36, 49, and 62 above respectively, and further in view of Junkin, U.S. Patent No. 6,493,717 (Junkin) and Harel, U.S. Patent No. 5,873,081 (Harel). In paragraph (13), the Office Action rejected claims 40, 53, and 66 under 35 U.S.C. §103(a) as being unpatentable over Callahan, Epperson, Posse, Junkin, and Harel as applied to claims 39, 52, and 65 above respectively, and further in view of Zhou et al., U.S. Patent No. 5,995,511 (Zhou). The Applicants respectfully traverse these rejections.

Claim 39 recites that the execution trace information includes a session identifier and a query ID, while claim 40 recites that the execution trace record includes an operator ID, a start timestamp, and a finish timestamp. Here again, the Office Action rejects these claims based upon references that at best, merely discloses a parameter, but does not disclose providing that parameter in anything analogous to an execution trace record.

As for a rationale to combine these references, the Office Action indicates:

"It would have been obvious to one of ordinary skill in the art at the time of the invention to combine Jurkin and with Callahan, Epperson, and Posse since Callahan, Epperson, Posse, and Jurkin teach the execution of tasks in the use of trace information, Callahan, Epperson, and Jurkin teach the use of networks and the use of routines and subroutines, Epperson and Jerkin teach the performing of query is, the use of databases, and the use of servers, posse and Jurkin teach the use of nodes."

The Applicants respectfully disagree that the foregoing rationale for combine these references comports with the requirements of 35 U.S.C. § 103. Even if the above statements were true in terms of what each reference teaches, it does not provide a rationale for combine or modify these references. Accordingly, the Applicants respectfully traverses the rejection of claims 39 and 40.

Claims 52 and 65 recite features analogous to those of claim 39, and are patentable on the same basis.

Claims 53 and 66 recite features analogous to those of claim 40, and are patentable on the same basis.

In paragraph (14), the Office Action rejected claims 41, 42, 54, 55, 67, and 68 under 35 U.S.C. §103(a) as being unpatentable over Callahan, Epperson, Posse, Junkin, Harel, and Zhou as applied to claims 40, 53, and 66 above, and further in view of Barnford et al., U.S. Patent No. 6,243,702 (Barnford). The Applicants respectfully traverse these rejections.

According to the Office Action, the Zhou reference teaches the is a start timestamps and finish timestamps. However, again, those timestamps are not used in anything analogous to an execution trace record, and no motivation to include those timestamps in execution trace record is provided. Hence, the Applicants respectfully traverse these rejections.

In paragraph (16), the Office Action rejected claims 44, 57, and 70 under 35 U.S.C. §103(a) as being unpatentable over Callahan, Epperson, and Posse as applied to claims 36, 49, and 62 above

respectively, and further in view of Naidu et al., U.S. Patent No. 5,752,002 (Naidu). The Applicants respectfully traverse these rejections.

According to the Office Action, Naidu teaches writing to different files at different times. Again, this teaching is out of context with anything analogous to an execution trace record, and no motivation is provided. Hence, the Applicants respectfully traverse these rejections.

In paragraph (17), the Office Action rejected claims 46, 48, 59, 61, 72, and 74 under 35 U.S.C. §103(a) as being unpatentable over Callahan, Epperson, and Posse as applied to claims 36, 49, and 62 above respectively, and further in view of Filepp et al., U.S. Patent No. 5,594,910 (Filepp). The Applicants respectfully traverse these rejections.

The Office Action indicates that Filepp teaches the use of presentation commands and timestamps, but, as before, not in the context of assembling execution trace information. Further, no mention is made of synchronizing execution trace record as according to the timestamp, as recited in claim 48. Accordingly, the Applicants respectfully traverse these rejections.

In paragraph (19), the Office Action rejected claims 47, 60, and 73 under 35 U.S.C. §103(a) as being unpatentable over Callahan, Epperson, and Posse as applied to claims 36, 49, and 62 above respectively, and further in view of Carey et al., U.S. Patent No. 6,285,997 (Carey). The Applicants respectfully traverse these rejections.

The Office Action indicates that the Carey reference teaches performing other routines while executing a query. However, the subject claims recite that execution trace information are generated in two separate entities (the query coordinator and the data server) while executing the query. Even if the Carey reference discloses performing other routines while executing a query, it does not disclose all of the features of the subject claims. Accordingly, the Applicants respectfully traverse these rejections.

## VI. Dependent Claims

Dependent claims 38-48, 51-61, and 64-74 incorporate the limitations of their related independent claims, and are therefore patentable on this basis. In addition, as described above, these claims recite novel elements even more remote from the cited references. Accordingly, the Applicants respectfully request that these claims be allowed as well.

**VII. Conclusion**

In view of the above, it is submitted that this application is now in good order for allowance and such allowance is respectfully solicited. Should the Examiner believe minor matters still remain that can be resolved in a telephone interview, the Examiner is urged to call Applicants' undersigned attorney.

Respectfully submitted,

GATES & COOPER LLP  
Attorneys for Applicant(s)

Howard Hughes Center  
6701 Center Drive West, Suite 1050  
Los Angeles, California 90045  
(310) 641-8797

Date: April 14, 2003

By: Victor G. Cooper  
Name: Victor G. Cooper  
Reg. No.: 39,641

VGC/sjm  
G&C 30145.344-US-01